

# Light and Widely Applicable MCMC: Approximate Bayesian Inference for Large Datasets

Florian Maire<sup>a,b,\*</sup>, Nial Friel<sup>a,b</sup>, Pierre Alquier<sup>c</sup>

<sup>a</sup>*School of Mathematical Sciences, University College Dublin, Ireland*

<sup>b</sup>*Insight Centre for Data Analytics, University College Dublin, Ireland*

<sup>c</sup>*ENSAE-CREST, Malakoff, France*

---

## Abstract

Light and Widely Applicable (LWA-) MCMC is a novel approximation of the Metropolis–Hastings kernel targeting a posterior distribution defined on a large number of observations. Inspired by Approximate Bayesian Computation, we design a Markov chain whose transition makes use of an *unknown* but *fixed*, fraction of the available data, where the random choice of sub-sample is guided by the fidelity of this sub-sample to the observed data, as measured by summary (or sufficient) statistics. LWA–MCMC is a generic and flexible approach, as illustrated by the diverse set of examples which we explore. In each case LWA–MCMC yields excellent performance and in some cases a dramatic improvement compared to existing methodologies.

*Keywords:* Approximate Bayesian Computation, Bayesian inference, Big-data, Fixed computational budget, noisy Markov chain Monte Carlo algorithm

*2010 MSC:* 65C40, 65C60, 62F15

---

## 1. Introduction

The development of statistical methodology which scales to large datasets represents a significant research frontier in modern statistics. This paper presents a generic and flexible approach to directly address this challenge. Given a set of observed data  $(Y_1, \dots, Y_N)$ , a specified prior distribution  $p$  and a likelihood function  $f$ , estimating parameters  $\theta \in \Theta$  of the model proceeds via exploration of the posterior distribution  $\pi$  defined on  $(\Theta, \mathcal{B}(\Theta))$  by

$$\pi(d\theta \mid Y_1, \dots, Y_N) \propto f(Y_1, \dots, Y_N \mid \theta)p(d\theta). \quad (1)$$

Stochastic computation methods such as Monte Carlo methods allow one to estimate characteristics of  $\pi$ . In Bayesian inference, Markov chain Monte Carlo (MCMC) methods remain the most widely used strategy. Paradoxically, improvements in data acquisition technologies together with increased storage capacities, present a new challenge for these methods. Indeed, the size of the data set  $N$  (along with the dimension of each observation) can become so large, that even a routine likelihood evaluation is made prohibitively computationally intensive. As a consequence, methods such as the Metropolis–Hastings

---

\*Corresponding author

Email address: `florian.maire@ucd.ie` (Florian Maire)

sampler cannot be realistically considered. This issue has recently generated a lot of research activity.

While some authors have designed *exact* algorithms that tend to match the theoretical requirements of usual MCMC methods, others have considered the possibility of *approximate*, or *noisy*, methods while still trying to derive some quantitative error bound on the resulting approximation scheme. In this context, we refer to as *exact* any method that produces –possibly dependent– samples from  $\pi$ , as opposed to *noisy*, which samples from an approximation of  $\pi$ . In Pseudo-marginal algorithms (Andrieu and Roberts, 2009), the evaluation of the likelihood function is substituted by an unbiased and positive estimator, while still preserving the stationary distribution  $\pi$  and thus being *exact*. Although theoretically appealing (Andrieu and Vihola, 2015), finding an unbiased and positive estimator of the likelihood turns out to be a challenging problem in itself. Two other more recent *exact* approaches (Banterle et al., 2014; Maclaurin and Adams, 2014) overcome the need for an unbiased and positive estimator of the likelihood. However, the method proposed in Maclaurin and Adams (2014) requires the specification of a lower bound of the likelihood – which is a strong assumption especially for high-dimensional problems. A poor lower bound function yields a method whose computational complexity can even exceed that of a M–H sampler. In Banterle et al. (2014), the authors suggest to break the usual M–H ratio into  $N$  independent decision steps (each corresponding to a factor involving the likelihood of a datum) such that a proposal is globally rejected as soon as it is rejected by an elementary step. This *sticky* version of the M–H sampler is nevertheless shown to be  $\pi$ -stationary but yields an higher asymptotic variance by a straightforward Peskun comparison argument with the standard M–H kernel (Tierney, 1998).

An alternative solution consists of finding an approximation of the M–H transition kernel, that would emulate the outcome of the accept/reject step without having to compute the likelihood ratio in the M–H transition kernel. In Bayesian settings where the observed data are independent and identically distributed (*i.i.d.*), the likelihood is expressed as

$$f(Y_1, \dots, Y_N | \theta) \propto \prod_{k=1}^N f(Y_k | \theta).$$

The issue of the reliability of making a decision to accept or reject a move based only a subset of these factors has been recently addressed (Korattikara et al., 2014; Bardenet et al., 2014). In these papers, the usual M–H acceptance decision, which can be rewritten as

$$\frac{1}{N} \log \left( U \frac{p(\theta)q(\theta, \theta')}{p(\theta')q(\theta', \theta)} \right) \leq \frac{1}{N} \sum_{k=1}^N \log \frac{f(Y_k | \theta')}{f(Y_k | \theta)}, \quad U \sim \mathcal{U}(0, 1) \quad (2)$$

is made using a Monte Carlo approximation

$$\rho_n = \frac{1}{n} \sum_{k=1}^n \log \frac{f(Y_{u_k} | \theta')}{f(Y_{u_k} | \theta)}, \quad (u_1, \dots, u_n) \in \{1, \dots, N\}^n, \quad (u_i \neq u_j)$$

of the right hand side of (2). Both methods proposed in Korattikara et al. (2014) and Bardenet et al. (2014) share the same principle that

- (a) a decision can be made with a certain *level of confidence* as soon as  $\rho_n$  becomes *sufficiently far* from the left-hand side of (2),

(b) and if this condition is not reached,  $n$  is increased.

They nevertheless differ by the way the level of confidence is derived and by the theoretical arguments motivating the approximation. However, practically, as observed in some examples presented in Korattikara et al. (2014) and Bardenet et al. (2014), both methods tend to draw a significant portion of the data (*i.e.*  $n \rightarrow N$ ), in order to reach the confidence interval when the Markov chain gets closer to stationarity. Finally, note that *noisy* algorithms may retain some theoretical appealing feature. This was developed in Alquier et al. (2014), where the authors show that using an approximation of an unavailable exact transition kernel can still provide ergodic Markov chains. In particular, this yields a class of general noisy M–H algorithms, extending the pseudo-marginal approach (Andrieu and Roberts, 2009), while relaxing the unbiased estimator assumption.

In this paper, we propose *Light and Widely Applicable MCMC (LWA–MCMC)*, a novel methodology which aims to make the best use of a computational resource available for a given computational run-time, while still preserving the celebrated simplicity of the standard M–H sampler. Our approach designs a Markov chain on an extended state space whose marginal in  $\theta$  targets an approximation of  $\pi$ . As a result, our algorithm can be cast as a *noisy* MCMC method. At each transition of the Markov chain, a new candidate is proposed and accepted/rejected through a probability that only uses a fraction,  $n/N$ , of the available data which is by construction –and contrary to Maclaurin and Adams (2014), Korattikara et al. (2014) and Bardenet et al. (2014)– held constant throughout the algorithm. Moreover, unlike most of the papers mentioned before, LWA–MCMC can be applied to virtually any model (involving *i.i.d.* data or not), as it does not require any assumption on the likelihood function nor on the prior distribution.

The original target  $\pi$  is extended to model a joint distribution between the parameter of interest  $\theta \in \Theta$  and an auxiliary  $N$ -dimensional boolean vector identifying the data involved in the subset. Each possible subset of data of size  $n$  is weighted according to a *similarity measure* with respect to the full set of data, in the spirit of the Approximate Bayesian Computation (ABC) (see *e.g.* Marin et al., 2012). In the special case of *i.i.d.* realizations from an exponential model, we prove that when the similarity measure is identical the sufficient statistics, this yields an optimal approximation, in the sense of minimizing an upper bound of the Kullback-Leibler (KL) divergence between  $\pi$  and the marginal target of our method.

Our main finding is two-fold:

- for a fixed computational budget, our method can achieve a better bias/variance tradeoff compared to a standard Metropolis–Hastings algorithm and the method developed in Bardenet et al. (2014)
- we observe in different scenarios the existence of a tradeoff between the *quality* and the *size*  $n$  of the batch of the sub-sampled of data, highlighting an LWA–MCMC *optimal* setting.

We start in Section 2 with a striking real data example which we hope will help the reader to understand the problem we address and motivate the solution we propose, without going into further technical details at this stage. In Section 3, we provide theoretical results concerning exponential-family models, which we illustrate through a probit example. This section allows us to justify our motivations supporting the LWA–MCMC general methodology developed in Section 4: we define the transition kernel on the extended state

space and show that it yields a Markov chain targeting, marginally, an approximation of  $\pi$ . Finally, in Section 5, our method is used to estimate parameters of a time-series model and to perform a binary classification task. In the latter example, we compare the performance of our algorithm with the SubLikelihood approach proposed in Bardenet et al. (2014).

## 2. An introductory example

We address the problem of estimating template shapes of handwritten digits from the MNIST database (<http://yann.lecun.com/exdb/mnist/>) by inferring a partially known deformable template model (Allasonnière et al., 2007). Here, a  $16 \times 16$  matrix represents a digit whose conditional distribution given its class  $(0, 1, \dots, 9)$  corresponds to a random deformation of the template shape, parameterized by a  $d = 256$  dimensional vector  $\theta$ , of the known digit. Assuming small deformations, we can rewrite the model as a standard regression problem:

$$\text{given } I_k = i, \quad Y_k = \phi(\theta_i) + \sigma^2 \epsilon_k \quad (3)$$

where  $I_k \in \{0, \dots, 9\}$  is the class of observation  $Y_k$  (regarded as a vector  $\mathbb{R}^{255}$ ),  $\phi : \mathbb{R}^{256} \rightarrow \mathbb{R}^{225}$  is some deterministic linear mapping,  $\sigma > 0$  is a variance parameter and  $\epsilon_j \sim \mathcal{N}(0_{225}, \text{Id}_{225})$  is additive noise. Given a set of  $N = 10,000$  labeled images and defining a prior distribution for  $\theta = \{\theta_1, \dots, \theta_9\}$ , one can estimate  $\theta$  through its posterior distribution, for example using a standard Metropolis–Hastings algorithm. However, for two main reasons, the efficiency of such an approach can be questioned: (i) computing the  $N$  likelihoods in the M–H ratio dramatically slows down each transition and (ii) the highly peaked posterior distribution hinders a quick exploration of the state space.

Based on these observations, our approach aims at working with different subsets of data which addresses issues (i) and (ii). At this stage, we do not provide precise details of the LWA–MCMC machinery nor its accuracy but we rather provide an insight of the rationale of this approach: we design a Markov chain whose transition kernel targets the posterior distribution of the parameter of interest  $\theta$  given a random subset of  $n$  data ( $n \ll N$ ). More specifically, we inject in the standard M–H transition a decision about *refreshing* the subset of data, which, as a result, will change randomly over time. In this example, we use the knowledge of the observation labels to promote subsets in which the proportion of each digit is balanced.

We considered only five digits,  $1, \dots, 5$  (for illustration purposes), subsets of size  $n = 100$  and a non-informative Gaussian prior for  $\theta$ , as specified in Allasonnière et al. (2007). Figure 1 indicates a striking advantage of our method compared to a standard M–H using the  $N = 10,000$  data. In this scenario, we allow a given computational budget (1 hour) for both methods and compare the estimation of the mean estimate of the two Markov chains. Qualitatively, the upper part of Figure 1 compares the estimated template shapes of the five digits at different time steps and shows that our method allows to extract template shapes much quicker than the standard M–H, while still reaching an apparent similar graphical quality asymptotically (after one hour). This fact is confirmed quantitatively, in the lower part of Figure 1, which plots, against time and for both methods, the Euclidean distance between the Markov chain mean estimate and the Maximum Likelihood Estimate  $(\theta_1^*, \dots, \theta_5^*)$  obtained using a stochastic EM (Allasonnière et al., 2007). More precisely,

we compare the real valued function  $\{d(t), t \in \mathbb{R}\}$  defined as

$$d(t) = \sum_{i=1}^5 \|\theta_i^* - \mu(\theta_{i,1:\kappa(t)})\|, \quad \text{where} \quad \begin{cases} \forall t \in \mathbb{R}, \kappa(t) = \max_{k \in \mathbb{N}} \{t \geq \tau_k\}, \\ \tau_k \text{ is the time at the end of the } k\text{-th iteration}, \\ \forall k \in \mathbb{N}, \mu(\theta_{i,1:k}) = (1/k) \sum_{j=1}^k \theta_{i,j}, \end{cases}$$

for the two Markov chains.

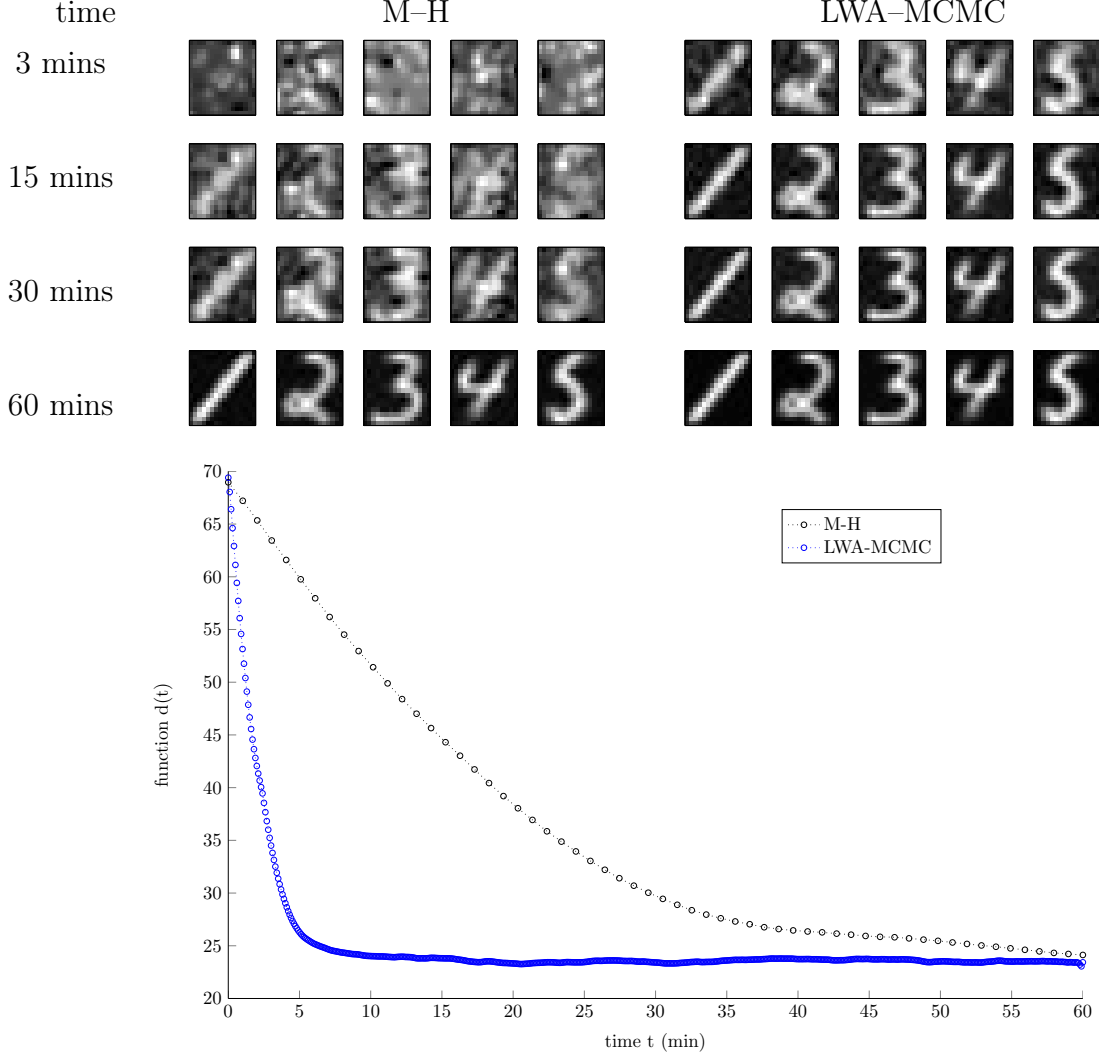


Figure 1: Efficiency of template estimation through M-H and LWA-MCMC.

LWA-MCMC provides very encouraging results for this real-data example. We formalize the method and sketch theoretical arguments in the next two sections.

### 3. Approximation of the posterior distribution in exponential models: an optimality result

In this section, we consider the case of  $N$  independent and identically distributed (*i.i.d.*) observations from an exponential model. Taking the posterior distribution given

all the  $N$  available data as the target distribution, which we call the *full posterior*, we consider the posterior distribution of the parameter of interest given *only* a subset of the  $N$  observations as an approximation of the full-posterior, which we call a *sub-posterior*. We investigate the influence of the choice of a subset of  $n$  data on this approximation. Proposition 1 shows the existence of an optimal set of possible subsets of size  $n$  with respect to the Kullback-Leibler (KL) divergence between the full posterior and the sub-posterior. This result will be used in the next section to design and justify the LWA-MCMC methodology, extending this approach to non-*i.i.d.* observations from general likelihood models.

### 3.1. Notation

Let  $(Y_1, \dots, Y_N) \in \mathbf{Y}^N$  be a set of *i.i.d.* observed data ( $\mathbf{Y} \subseteq \mathbb{R}^m$ ,  $m > 0$ ) and define

- $Y_{i:j} = (Y_i, \dots, Y_j)$  if  $1 \leq i \leq j \leq N$  with the convention that  $Y_{i:j} = \{\emptyset\}$ , otherwise.
- $Y_{-k} = (Y_1, \dots, Y_{k-1}, Y_{k+1}, \dots, Y_N)$  for all  $k \in \{1, \dots, N\}$ .
- $Y_U = \{Y_k, k \in U\}$ , where  $U \subseteq \{1, \dots, N\}$ .

In this section, we assume that the likelihood model  $f$  belongs to the exponential family and is fully specified by a vector of parameters  $\theta \in \Theta$ , ( $\Theta \subseteq \mathbb{R}^d$ ,  $d > 0$ ) and a sufficient statistic mapping  $S : \mathbf{Y} \rightarrow \mathbf{S}$  ( $\mathbf{S} \subseteq \mathbb{R}^s$ ,  $s > 0$ ) such that

$$f(y | \theta) = \exp \langle g(\theta), S(y) \rangle / L(\theta)$$

is the density of the likelihood distribution with respect to the Lebesgue measure  $\lambda(dy)$ . Here, the symbol  $\langle \cdot, \cdot \rangle$  denotes the canonical inner product in  $\mathbf{S}$ ,  $g : \Theta \rightarrow \mathbf{S}$  is a model-specific mapping and  $L(\theta)$  is the likelihood normalizing constant.

The posterior distribution  $\pi$  is defined on the measurable space  $(\Theta, \vartheta)$  by its density function

$$\pi(\theta | Y_{1:N}) = p(\theta) \prod_{k=1}^N f(Y_k | \theta) / \int p(d\theta) \prod_{k=1}^N f(Y_k | \theta), \quad (4)$$

with respect to the Lebesgue measure on  $(\Theta, \vartheta)$ .  $p$  is a prior distribution defined on  $(\Theta, \vartheta)$  and with some abuse of notation,  $p$  denotes also the probability density function (pdf) on  $\Theta$  accordingly ( $p(d\theta) = p(\theta)\lambda(d\theta)$ ).

For all  $n \leq N$ , we define  $\mathbf{U}_n$  as the set of the possible combinations of  $n$  different integer numbers less than or equal to  $N$  and  $\mathcal{U}_n = 2^{\mathbf{U}_n}$  as the powerset of  $\mathbf{U}_n$ . Finally, let  $\mathbf{U} = \bigcup_{n \leq N} \mathbf{U}_n$  and  $\mathcal{U} = 2^{\mathbf{U}}$ . In the sequel, we set  $n$  as a constant and wish to compare the full-posterior distribution (4) with any of the sub-posterior distributions from the family  $\mathbf{F}_n = \{\tilde{\pi}(\cdot | Y_{1:N}, U_n), U_n \in \mathbf{U}_n\}$ , where for all  $U_n \in \mathbf{U}_n$ , we have defined

$$\tilde{\pi}(\theta | Y_{1:N}, U_n) = \tilde{\pi}(\theta | Y_{U_n}) = p(\theta) \prod_{k \in U_n} f(Y_k | \theta) / \int p(d\theta) \prod_{k \in U_n} f(Y_k | \theta). \quad (5)$$

### 3.2. Optimal subsets for the Kullback-Leibler divergence between $\pi$ and $\tilde{\pi}$

Recall that for two measures  $\pi$  and  $\tilde{\pi}$  defined on the same measurable space  $(\Theta, \vartheta)$ , the Kullback-Leibler (KL) divergence between  $\pi$  and  $\tilde{\pi}$  is defined as:

$$\text{KL}(\pi \| \tilde{\pi}) = \mathbb{E}_{\pi} \left\{ \log \left( \frac{\pi(\theta)}{\tilde{\pi}(\theta)} \right) \right\}. \quad (6)$$

Although not a proper distance between probability measures defined on the same measurable space,  $\text{KL}(\pi \| \tilde{\pi})$  is used as a similarity criterion between  $\pi$  and  $\tilde{\pi}$ . It can be interpreted in information theory as a measure of the information lost when  $\tilde{\pi}$  is used to approximate  $\pi$ , which is our primary concern here. We now state the main result of this section.

**Proposition 1.** *Consider the KL divergence between  $\pi = \pi(\cdot | Y_{1:N})$  and  $\tilde{\pi}_U = \tilde{\pi}(\cdot | Y_U) \in \mathbf{F}_n$ , we have:*

(i)

$$\text{KL}(\pi \| \tilde{\pi}_U) \leq \Psi(n, N, Y_{1:N}) + B(U, n, Y_{1:N}) \quad (7)$$

where  $\Psi$  is a deterministic function independent of  $U$  and  $B$  is a positive function.

(ii) If the set of optimal subsets  $\mathbf{U}_n^* \subset \mathbf{U}_n$  defined by

$$\mathbf{U}_n^* := \left\{ U \in \mathbf{U}_n, \quad \frac{1}{N} \sum_{k=1}^N S(Y_k) = \frac{1}{n} \sum_{k \in U} S(Y_k) \right\} \quad (8)$$

is non-empty, then for any  $U \in \mathbf{U}_n^*$ ,  $B(U, n, Y_{1:N}) = 0$ , hence yielding an optimal upper-bound.

(iii) For a given subset  $U_0 \in \mathbf{U}_n$  such that for all  $U \in \mathbf{U}_n$

$$\left\| \frac{1}{N} \sum_{k=1}^N S(Y_k) - \frac{1}{n} \sum_{k \in U_0} S(Y_k) \right\| \leq \left\| \frac{1}{N} \sum_{k=1}^N S(Y_k) - \frac{1}{n} \sum_{k \in U} S(Y_k) \right\|, \quad (9)$$

then we have

$$B(U_0, n, Y_{1:N}) \leq B(U, n, Y_{1:N}). \quad (10)$$

In other words, the sub-posterior distributions in  $\mathbf{F}_n$  with subsets having the same sufficient statistics on average as for the full dataset, will achieve an optimal approximation (with respect to upper-bounding the KL divergence). Moreover (9) defines an order on  $\mathbf{F}_n$  which implies an order on their relative KL divergence upper-bound (10). The proof is outlined in Appendix Appendix A.

### 3.3. Illustration with a probit model: effect of choice of sub-sample

We consider a pedagogical example, based on a probit model, to illustrate Proposition 1. A probit model is used in regression problems in which a binary variable  $Y_k \in \{0, 1\}$  is observed through the following sequence of independent random experiments, defined for all  $k \in \{1, \dots, N\}$  as:

(i) Draw  $X_k \sim \mathcal{N}(\theta^*, \gamma^2)$

$r(U)$	$\text{KL}(\pi \  \tilde{\pi}_U) / \text{KL}(\pi \  \tilde{\pi}_{U^*})$
0.01	1.13
0.04	1.39
0.07	1.87
0.1	2.76

Table 1: Comparison of the KL divergence between the full-posterior and the optimal sub-posterior ( $r(U) = 0$ ) with the KL distance between the full-posterior and other sub-posterior distribution.

(ii) Set  $Y_k$  as follows

$$Y_k = \begin{cases} 1, & \text{if } X_k > 0, \\ 0, & \text{otherwise.} \end{cases} \quad (11)$$

Observing a large number of realizations  $Y_1, \dots, Y_N$ , we aim to estimate the posterior distribution of  $\theta$ . In practice, one also estimates  $\gamma$  but for illustration purpose, this parameter is considered as known here. The likelihood function can be expressed as

$$f(Y_k | \theta) = \alpha(\theta)^{Y_k} (1 - \alpha(\theta))^{(1-Y_k)} = (1 - \alpha(\theta)) \left( \frac{\alpha(\theta)}{1 - \alpha(\theta)} \right)^{Y_k}, \quad (12)$$

where  $\alpha(\theta) = \int_0^\infty (2\pi\gamma^2)^{-1/2} \exp\{(1/2\gamma^2)(t-\theta^*)^2\} dt$  and clearly belongs to the exponential family. The full-posterior distribution can be written as

$$\pi(\theta | Y_{1:N}) \propto p(\theta) (1 - \alpha(\theta))^N \left( \frac{\alpha(\theta)}{1 - \alpha(\theta)} \right)^{\sum_{k=1}^N Y_k},$$

where  $p$  is the prior density, which we assume to be non informative ( $p(\theta) = \mathcal{N}(a, b^2)$ ). In this example, the mapping  $\theta \rightarrow \pi(\theta | Y_{1:N})$  can be easily estimated for any  $\theta \in \Theta$ , even when  $N$  is extremely large, as it only requires to sum over all the binary variables  $Y_1, \dots, Y_N$ . As a consequence, samples from the full-posterior distribution  $\pi(\cdot | Y_{1:N})$  can be routinely obtained by a standard M-H algorithm and similarly for any sub-posterior distributions  $\tilde{\pi}(\cdot | Y_U) \in \mathbf{F}_n$ .

We present in Figure 2 some inference results for the full-posterior and several sub-posterior distributions (with different  $n$  and different values of sufficient statistics) obtained with parameters  $(\theta^*, \gamma) = (1, 1)$  and  $N = 10,000$  simulated data. In the upper plot, we hold  $n = 100$  constant and compare several sub-posterior distributions given subsets of data  $U \in \mathbf{U}_{100}$  having different matches with the full data sufficient statistics  $r(U) = |\bar{S}(U) - \bar{S}_N|$ , where  $\bar{S}(U) = n^{-1} \sum_{k \in U} S(Y_k)$  and  $\bar{S}_N = N^{-1} \sum_{k=1}^N S(Y_k)$ . In this probit model,  $S$  is simply the identity function, implying that  $r(U)$  monitors the different proportion of 1 and 0's between the full dataset and in the subset  $U$ . This plot, as well as the quantitative result of Table 1 providing the KL divergence between the full-posterior and these different sub-posterior distributions are consistent with the statement of Proposition 1: when learning from a subset of  $n$  data, one should work with a subset  $U$  featuring a perfect match with the full dataset, *i.e*  $r(U) = 0$ , to achieve an *optimal* approximation of  $\pi$ . Note that the sub-posterior distributions only vary through their sufficient statistics, so obviously, in exponential models, the choice of the *optimal* subset is not unique, as  $\mathbf{U}_n^*$  is not restricted to one element.



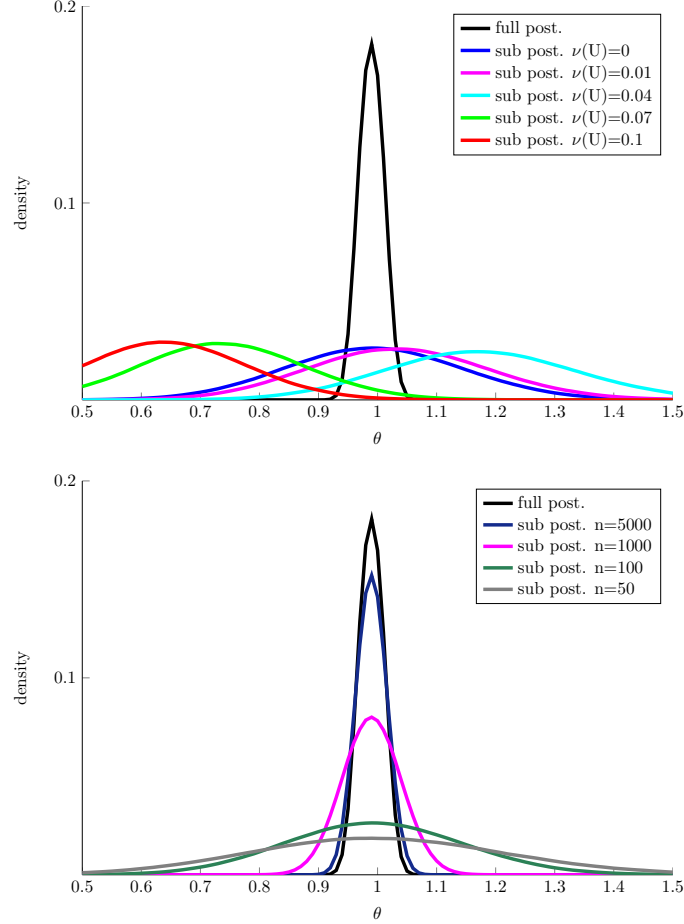


Figure 2: Influence of the parameter  $U \in \mathbf{U}_{100}$  on the sub-posterior distribution  $\pi(\cdot | Y_U)$  and comparison with  $\pi$  (top) – Influence of the subsets size  $n$  on the sub-posterior distribution  $\pi(\cdot | Y_U)$ ,  $U \in \mathbf{U}_n^*$  (bottom).

The lower plot of Figure 2 compares the influence of  $n$  on the *optimal* sub-posterior distributions  $\tilde{\pi}(\cdot | Y_U)$ ,  $U \in \mathbf{U}_n^*$  for  $n \in \{50, 100, 1000, 5000\}$ . As expected, while the variance becomes wider as  $n$  decreases, the expectation remains relatively constant.

**Remark 1.** *The results of the lower plot of Figure 2 may in some ways be related with the local asymptotic normality (lan) theorem: Bernstein von Mises theorem (see e.g. Van der Vaart, 2000) states that under some mild assumptions about the likelihood and the prior distribution, the posterior distribution is asymptotically Gaussian (in  $n$ ) in the case where the "true" parameter is an interior point of the parameter space. More precisely, the mean of the Gaussian distribution corresponds to the maximum likelihood estimate of the observed data, while the covariance is given by  $H^{-1}/n$ , where  $H$  is the Hessian matrix at the "true" parameter value. Even though for low values of  $n$ , such an asymptotic result does not hold, it is nevertheless consistent with observing in Figure 2 that the variance varies (in function of  $n$ ) much more than the mean of those distributions.*

**Remark 2.** *At this stage, one might wonder why so much effort has been put to overcome the problem of sampling from a posterior distribution given a huge amount of data, while using a local asymptotic normality theorem would virtually allow to solve this problem by sampling from a Gaussian distribution. Two main arguments actually prevent one to make use of such a result:*

- *estimation of the coefficients of the asymptotic Gaussian distribution is non-trivial, as one needs to invert the Hessian matrix at  $\theta^*$ , which is typically unknown;*
- *local asymptotic normality theorems only hold under restrictive assumptions, for example, when the observations are i.i.d.*

On the basis of our analysis conducted in the case of *i.i.d.* realizations from an exponential-family model, both at theoretical and experimental levels, it seems reasonable to consider estimating the Maximum A Posteriori parameter based on a subset of data. At first glance, when one aims to estimate  $\theta$  using a Markov chain targeting a sub-posterior, even optimal,  $\pi(\cdot | Y_U)$ , ( $U \in \mathbf{U}_n^*$ ) instead of the full-posterior may lead to a worse efficiency, iteration wise. However, assuming that the computational complexity of a Markov chain targeting the full posterior is prohibitively intensive, one may consider a Markov chain targeting an optimal sub-posterior as a realistic alternative: more Markov chain iterations would be required but at a known and affordable computational cost. This yields a trade on the subset size  $n$  which will allow lighter transitions at the price of a loss in variance.

#### 4. Light and Widely Applicable MCMC: the general methodology

In this section, we do not assume any particular correlation pattern for the sequence of observations, nor any specific likelihood model and simply write the posterior distribution  $\pi$  as

$$\pi(d\theta | Y_1, \dots, Y_N) \propto p(d\theta) f(Y_{1:N} | \theta). \quad (13)$$

The Light and Widely Applicable MCMC (LWA-MCMC) methodology that we describe now can be regarded as a generalization of the sub-posterior inference detailed in the previous section to non-exponential-family models with possibly dependent observations.

##### 4.1. Motivation of our approach

Here, we do not assume the existence of a sufficient statistic mapping for the model under consideration. Thus, in order to allow comparison between different subsets of data, we introduce an *artificial* summary statistic mapping  $S_n : \mathbf{Y}_n \rightarrow \mathbf{S}$  ( $n \leq N$ ), where  $\mathbf{S} \subseteq \mathbb{R}^s$ . The choice of the summary statistics  $S_n$  is problem specific and is meant to be the counterpart of the sufficient statistic mapping for general models (hence sharing, slightly abusively, the same notation). Intuitively, a good choice of  $S_n$  would capture most of the statistical features of  $Y_U$  ( $U \in \mathbf{U}_n$ ). In an attempt to derive a similar analysis to the exponential model case (Section 3) and to reach an optimal setup in line with Proposition 1, we want to focus inference on those subsets whose summary statistics vector is close to that of the full dataset. In this context a subset  $Y_U$  ( $U \in \mathbf{U}_n$ ) is said to be more *representative of the full dataset* than another  $Y_{U'}$  ( $U' \in \mathbf{U}_n$ ), if  $\|\tilde{S}_n(Y_U) - \tilde{S}_N(Y_{1:N})\| \leq \|\tilde{S}_n(Y_{U'}) - \tilde{S}_N(Y_{1:N})\|$ , where we have set  $\tilde{S}_n(Y_{U'}) = S_n(Y_{U'})/n$  as the *normalized* summary statistics and  $\|\cdot\|$  as the Euclidean distance on  $\mathbf{S}$ . Since the question of specifying summary statistics also arises in ABC, one can take advantage of the abundant ABC literature on this topic to find some examples of summary statistics for usual likelihood models (see *e.g.* Nunes and Balding, 2010; Csilléry et al., 2010; Marin et al., 2012; Fearnhead and Prangle, 2012).

We formalise this idea by assigning a weight to any possible subset of data  $U \in \mathbf{U}_n$ . Let for any  $n \leq N$  and  $\epsilon > 0$ ,  $\nu_{n,\epsilon}$  be the distribution defined on the discrete state space  $(\mathbf{U}_n, \mathcal{U}_n)$  whose density with respect to the counting measure is:

$$\forall U \in \mathbf{U}_n, \quad \nu_{n,\epsilon}(U) = \Phi \left( -\frac{\|\bar{S}_n(Y_U) - \bar{s}_N\|}{\epsilon} \right) \Bigg/ \sum_{U' \in \mathbf{U}_n} \Phi \left( -\frac{\|\bar{S}_n(Y_{U'}) - \bar{s}_N\|}{\epsilon} \right). \quad (14)$$

Here  $\Phi : \mathbb{R} \rightarrow \mathbb{R}^+$  is a kernel function,  $\bar{s}_N = \bar{S}_N(Y_{1:N}) \in \mathbf{S}$  and the dependence of  $\nu_{n,\epsilon}(U)$  on  $Y_{1:N}$  is implicit. The parameter  $\epsilon$  allows to control the influence of the representativeness of a subset  $U \in \mathbf{U}_n$  on its overall weight  $\nu_{n,\epsilon}(U)$ : if  $\epsilon \gg 1$ , the weights tend to be uniform, whereas if  $\epsilon \ll 1$  the weights tend to highlight the representativeness of the subsets. As a consequence,  $\epsilon$  is a tuning parameter whose impact on the inference is significant, as we will see in Section 5. The kernel  $\Phi$  allows to smooth the weights and offers protection against the possibly unbounded weights derived from Euclidean distances (situations which typically arise when weights are proportional to  $\|\bar{S}_n(U) - \bar{s}_N\|^{-1}$  and  $\mathbf{U}_n^*(8)$  is not the empty set). In this paper, we let  $\Phi$  be the Gaussian kernel.

**Remark 3.** *Note that because the statistics used to assess the representativeness of a subset w.r.t. the full dataset are only summary and not sufficient, two subsets  $(U, U') \in \mathbf{U}_n^2$  such that  $\nu_{n,\epsilon}(U) = \nu_{n,\epsilon}(U')$  might yield two different sub-posteriors. As a consequence, should a unique optimal subset  $U^* = \arg \max_{U \in \mathbf{U}_n} \nu_{n,\epsilon}(U)$  be available, inferring the full-posterior through this corresponding optimal sub-posterior is likely to provide an unbalanced learning as most of the data would be simply ignored. Alternatively, when learning from a set of good subsets, where goodness is measured by  $\nu_{n,\epsilon}$ , one can expect that each sub-posterior involved in the process will act complementarily to improve the approximation of  $\pi$ .*

At this stage, two main questions need to be addressed:

- (i) how can a set of *good* subsets be determined? Indeed, the dimension of  $\mathbf{U}_n$ ,  $|\mathbf{U}_n| = \binom{N}{n}$  is prohibitively too large to estimate the normalizing constant in (14) and therefore we reasonably assume in the following that the set of weights  $\{\nu_{n,\epsilon}(U), U \in \mathbf{U}_n\}$  is unknown.
- (ii) how to design an inference scheme that would make use of these subsets, without resorting to a set of independent Markov chains each targeting through a standard M–H transition kernel a good sub-posterior, which could become even more demanding than a standard M–H targeting  $\pi$ ?

Because we assume that  $\{\nu_{n,\epsilon}(U), U \in \mathbf{U}_n\}$  is unknown, we consider that the subsets  $U_1, U_2, \dots (U_i \in \mathbf{U}_n)$  involved in the learning setup are latent variables of the model. We define the data-augmented distribution  $\tilde{\pi}_{n,\epsilon}$  for any  $A \in \mathcal{V}$  and  $U \in \mathbf{U}_n$  by

$$\tilde{\pi}_{n,\epsilon}(A, U | Y_{1:N}) = \nu_{n,\epsilon}(U) \pi(A | Y_U), \quad (15)$$

where

$$\pi(d\theta | Y_U) = p(\theta) f(Y_U | \theta) \lambda(d\theta) \Bigg/ \int p(\theta') f(Y_U | \theta') \lambda(d\theta')$$

is a sub-posterior distribution.

With some abuse of notation, we write  $\tilde{\pi}_{n,\epsilon}$  also for the density of the data-augmented distribution with respect to the product measure  $\lambda(d\theta)\mathbb{1}_{\mathcal{U}_n}(U)$ . Integrating out the subset variable yields the marginal of interest whose density w.r.t.  $\lambda$  is expressed as:

$$\tilde{\pi}_{n,\epsilon}^*(\theta | Y_{1:N}) = \sum_{U \in \mathcal{U}_n} \nu_{n,\epsilon}(U) \pi(\theta | Y_U). \quad (16)$$

In this way, the marginal distribution  $\tilde{\pi}_{n,\epsilon}^*$  defines an approximation of  $\pi$ . Straightforwardly, as  $n \rightarrow N$ , our approximation  $\tilde{\pi}_{n,\epsilon}^*$  converges to  $\pi$ .

**Remark 4.** *The parameters  $n$  and  $\epsilon$  of the joint distribution  $(\theta, U)$ , have complementary roles in the approximation of  $\pi$ . While  $n \ll N$  lightens the full-posterior distribution,  $\epsilon$  allows to make up for the approximation by attaching bigger weights to representative subsets in the mixture distribution  $\tilde{\pi}_{n,\epsilon}$  (15). Setting  $\epsilon$  implies a tradeoff between:*

- $\epsilon \gg 1$  and  $\nu_{n,\epsilon}$  is a flat distribution on  $U$  and all subsets have the same weight in the mixture
- $\epsilon \ll 1$  and  $\tilde{\pi}_{n,\epsilon}$  has most of its probability mass on  $\Theta \times \mathcal{U}_n^*$

In particular, the latter scenario yields an approximation of the full-posterior based only on one or a few subsets. Depending on the relevance of the summary statistics, the choice  $\epsilon \ll 1$  may provide a disastrous setup. The simulations detailed in Section 5 clearly show the existence of an optimal  $\epsilon$  for a given  $n$  and which also depends on the likelihood model considered.

#### 4.2. Formal description of the LWA-MCMC algorithm

In this section, we specify LWA-MCMC which allows to sample from an approximation of  $\pi$ . This relies on a Markov chain whose transition kernel achieves our main target, namely having a bounded computational complexity, which can be controlled through the parameter  $n$ .

##### 4.2.1. LWA-MCMC

LWA-MCMC makes use of a Markov transition kernel operating on an extended state space  $(\Theta \times \mathcal{U}_n, \vartheta \otimes \mathcal{U}_n)$ . Algorithm 1 depicts a Markov transition of the LWA-MCMC algorithm. This embeds two successive decisions: the first one allows to *refresh* the subset variable  $U$  while the second one updates the parameter  $\theta$ .  $R(U, \cdot)$  and  $Q(\theta, \cdot)$  are proposal kernels respectively on  $\mathcal{U} \times \mathcal{U}$  and  $\Theta \times \mathcal{B}(\Theta)$ .




Figure 3: Intertwined structure of the LWA-MCMC Markov chain.

The dependence structure within the sequence of variables  $\{\theta_k, k \in \mathbb{N}\}$  and  $\{U_k, k \in \mathbb{N}\}$  produced by LWA-MCMC is displayed in Figure 3. This yields an *intertwined* design

---

**Algorithm 1** LWA–MCMC transition  $(\theta_k, U_k) \rightarrow (\theta_{k+1}, U_{k+1})$ 


---

- 1: **Input:** current state  $(\theta_k, U_k)$  and summary statistics of current subset  $S_k = \bar{S}_n(Y_{U_k})$
- 2: propose to refresh the subset  $U' \sim R(U_k, \cdot)$
- 3: compute the summary statistics  $S' = \bar{S}_n(Y_{U'})$
- 4: set  $U_{k+1} = U'$  with probability

$$\rho_{n,\epsilon}(U_k, U') = 1 \wedge \frac{R(U', U_k) \Phi(-\|S' - \bar{S}_N\|/\epsilon)}{R(U_k, U') \Phi(-\|S_k - \bar{S}_N\|/\epsilon)}, \quad (17)$$

- 5: leave  $U_{k+1} = U_k$  otherwise
- 6: **if** the subset has not been refreshed **then**
- 7:     propose a change of parameter  $\theta' \sim Q(\theta_k; \cdot)$
- 8:     set  $\theta_{k+1} = \theta'$  with probability

$$\hat{\alpha}(\theta_k, \theta' | U_k) = 1 \wedge \frac{\pi(\theta' | Y_{U_k}) Q(\theta'; \theta_k)}{\pi(\theta_k | Y_{U_k}) Q(\theta_k; \theta')}, \quad (18)$$

- 9:     leave  $\theta_{k+1} = \theta_k$  otherwise
  - 10: **else**
  - 11:     set  $\theta_{k+1} \sim T_{U_{k+1}}(\theta_k; \cdot)$  where  $T_{U_{k+1}}(\theta_k; \cdot)$  is the M–H transition kernel corresponding to steps 7, 8 and 9 iterated  $L$  times ( $L > 1$ )
  - 12: **end if**
  - 13: **return:**  $(\theta_{k+1}, U_{k+1})$  and  $S'$  if  $U_{k+1} \neq U_k$  and  $S_k$  otherwise
- 

in which  $\{U_k, k \in \mathbb{N}\}$  is itself a Markov chain targeting  $\nu_{n,\epsilon}$ , such that  $U_k$  is independent from the past sequence of parameters  $(\theta_1, \theta_2, \dots, \theta_k)$ . This is an important feature of our proposal as allowing  $U_{k+1}$  to depend on  $\theta_k$  could potentially lead to a local optimisation of the parameter  $\theta$  with respect to a given subset  $Y_U$ . Indeed, a parameter  $\theta_k$  may happen to fit particularly well the subset of data  $Y_{U_k}$  and therefore the  $U$  component of the chain may get stuck at  $U_k$  for a large number of iterations. We claim that LWA–MCMC allows to move freely throughout  $\mathbf{U}_n$  at a limited computational cost, a fact which is empirically confirmed by the simulations. Moreover, the two distinctive accept/reject steps for  $U$  and  $\theta$  allows the parameter  $\theta$  not to be hampered by a lack of move in the  $U$  direction.

To summarize, the LWA–MCMC sampling scheme is appealing as it allows:

- (i) to update the subset  $U_k$  and the parameter  $\theta_k$  through two different decisions,
- (ii) to make the subset update independent of  $\theta_k$ ,
- (iii) to control the computational complexity of a transition,
- (iv) to be applied in any inference problem where the likelihood function is tractable.

The following remark shows that standard Markov chain methods cannot be used to sample from  $\tilde{\pi}_{n,\epsilon}$ , hence justifying the sampling machinery of LWA–MCMC from another perspective.

**Remark 5.** *A block update M–H or a Metropolis-within-Gibbs algorithm targeting  $\tilde{\pi}_{n,\epsilon}$  cannot be implemented, as both involve at some stage an intractable acceptance ratio (the*

normalizing constant of the sub-posteriors does not cancel anymore because of the different subsets of data involved).

#### 4.2.2. Stability of LWA-MCMC

Let us assume that  $(\theta_k, U_k) \sim \tilde{\pi}_{n,\epsilon}$ . First, note that marginally

$$(\theta_k, U_k) \sim \tilde{\pi}_{n,\epsilon} \Rightarrow U_k \sim \nu_{n,\epsilon} \Rightarrow U_{k+1} \sim \nu_{n,\epsilon}, \quad (19)$$

where the later implication holds as  $\{U_k, k \in \mathbb{N}\}$  is in itself a M-H Markov chain with target distribution  $\nu_{n,\epsilon}$  (See Figure 3 and Algorithm 1 steps 2–5). On the one hand, given the event {no subset refresh},  $(\theta_{k+1}, U_{k+1}) \sim \tilde{\pi}_{n,\epsilon}$  simply derives from the fact that  $\theta_{k+1} | U_{k+1} = U_k \sim \pi(\cdot | Y_{U_k})$  which holds since steps 7–9 in Algorithm 1 consist in a standard Metropolis-within-Gibbs update for  $\theta$ . On the other hand, the event {subset refresh} disturbs the stationarity:  $(\theta_{k+1}, U_{k+1}) \not\sim \tilde{\pi}_{n,\epsilon}$ . Indeed, in order to remain invariant, the transition kernel should produce a sample  $\theta_{k+1}$  such that  $\theta_{k+1} | U_{k+1} \sim \pi(\cdot | U_{k+1})$ . However, this is not achievable in a single Markov transition since the target distribution  $\pi(\cdot | U_k)$  becomes instantaneously  $\pi(\cdot | U_{k+1})$ . Instead, the sample  $\theta_k$  can be regarded as the initial state of a Markov chain having  $\pi(\cdot | U_{k+1})$  as stationary distribution. Taking  $L \gg 1$  provides a state  $\theta_{k+1}$  which is approximately distributed under  $\pi(\cdot | U_{k+1})$ . Therefore, the LWA-MCMC transition kernel is not, *stricto-sensu*, stationary with respect to  $\tilde{\pi}_{n,\epsilon}$ , as when the subset is refreshed,  $\theta_{k+1}$  is a sample from the full conditional asymptotically in  $L$ . At this point we make the following remarks.

- Since the sequence of distributions  $\{\pi(\cdot | Y_{U_k}), k \in \mathbb{N}\}$  are likely to be close to each other, one can use a very limited number of intermediate steps, typically  $L = 1$  was used in all the simulations of Section 5. Indeed, depending on the number of data refreshed by  $R$ , two consecutive subsets  $Y_{U_k}$  and  $Y_{U_{k+1}}$  may only differ through very few observations, hence advocating setting  $L = 1$  does not adversely effect the convergence of the Markov chain, a statement which is supported by our experiments.
- Tuning the prior distribution with  $\epsilon \ll 1$  leads to a set of weights  $\{\nu_{n,\epsilon}(U), U \in \mathcal{U}_n\}$  with high discrepancy. Therefore, when the marginal chain  $\{U_k, k \in \mathbb{N}\}$  reaches stationarity, the subset samples yield similar summary statistics and makes the distributions  $\pi(\cdot | U_k)$  and  $\pi(\cdot | U_{k+1})$  even closer. In addition, this makes refreshing  $U$  an unlikely event and a transition is therefore most of the time valid.

Although, we do not carry out further theoretical analysis in this paper and leave the proof of stability of the marginal Markov chain  $\{\theta_k, k \in \mathbb{N}\}$  as the central question of a forthcoming paper, we outline two considerations that are likely to help proving the ergodicity of the LWA-MCMC Markov chain.

- Considering  $\tilde{\pi}_{n,\epsilon}$  as an intermediate target distribution

The ergodicity of the data-augmented Markov chain relies on the ability of the sampler to absorb minor target changes. To the best of our knowledge and probably as a result of a rather unusual MCMC development, there has been very little literature on the convergence of Markov chain toward time-evolving target distributions. However, we connect this question to the one addressed in Kuhn and Lavielle (2004) about the stability of a Markov chain targeting a sequence of distributions, each parameterized by a vector which is recursively updated through a Stochastic Approximation procedure (Robbins and Monro, 1951).

- Considering  $\pi$  as the ultimate target distribution

In another perspective, considering the marginal chain of interest  $\{\theta_k, k \in \mathbb{N}\}$  and regarding  $\{U_k, k \in \mathbb{N}\}$  as an auxiliary sequence of parameters, the acceptance probability  $\hat{\alpha}(\theta, \theta' | U)$  of the LWA–MCMC can be viewed as a noisy version of the M–H acceptance ratio, which is the acceptance probability involved in the M–H transition targeting the full-posterior  $\pi$ . This observation combined with the main result of Alquier et al. (2014) (see also Pillai and Smith, 2014) stating that a M–H transition kernel with a noisy acceptance probability nevertheless allows to obtain samples from a distribution whose total variation distance to  $\pi$  is bounded from above, provided that for all  $\theta_k \in \Theta$ ,  $U_{k+1} \in \mathcal{U}$  and  $\theta' \sim Q(\theta_k, \cdot)$ ,

$$\mathbb{E} |\hat{\alpha}(\theta_k, \theta' | U_{k+1}) - \alpha(\theta_k, \theta')| \leq \delta(\theta_k, \theta'),$$

where the expectation is taken under  $U_{k+1}$  and  $\delta : \Theta^2 \rightarrow \mathbb{R}^+$  is a deterministic function.

## 5. Illustrations of LWA–MCMC

We evaluate the efficiency of LWA–MCMC in two different applications: the first is posterior inference of a time series observed at  $N = 10^7$  contiguous time steps. This example is particularly relevant since the observations are non *i.i.d.*: as a result, this makes LWA–MCMC the only competing algorithm against M–H to infer such a model – the other solutions (Korattikara et al., 2014; Bardenet et al., 2014; Maclaurin and Adams, 2014; Banterle et al., 2014) being only suitable for *i.i.d.* data. The second task is a Gaussian binary classification problem based on  $N = 10^7$  data, which we use to compare LWA–MCMC with the M–H Sub Likelihood inference method proposed in Bardenet et al. (2014). Finally, we provide some additional details to the handwritten digit example outlined in Section 2.

### 5.1. Inference of an ARMA model

An ARMA(1,1) time series  $\{Y_k, k \leq N\}$  is defined recursively by:

$$\begin{cases} Y_0 \sim \mu, & Z_0 \sim \mathcal{N}(0, \sigma^2) \\ Y_k = \alpha Y_{k-1} + \beta Z_{k-1} + \gamma + Z_k, & Z_k \sim \mathcal{N}(0, \sigma^2), \quad \forall k \geq 1, \end{cases} \quad (20)$$

where  $\mu$  is some distribution on  $(Y, \mathcal{Y})$  and  $\theta = (\alpha, \beta, \gamma) \in \mathbb{R}^3$ . The likelihood of a trajectory can be written

$$f(Y_{0:N} | \theta) = \mu(Y_0) \prod_{k=1}^N g_k(Y_k | Y_{0:k-1}, \theta), \quad (21)$$

such that for all  $k \geq 1$ ,

$$g(Y_k | Y_{0:k-1}, \theta) = \Psi(Y_k; \alpha Y_{k-1} + \beta h_k(Y_{1:k-1}) + \gamma, \sigma^2) \quad (22)$$

where we have defined  $x \rightarrow \Psi(x; m, v)$  as the pdf of the univariate Gaussian distribution with mean  $m$  and variance  $v$  and  $\{h_k, k \geq 1\}$  is a set of known deterministic mappings.

The purpose of this example is to infer the posterior distribution  $\pi(\theta | Y_{1:N})$  using LWA-MCMC. We sampled a time series  $\{Y_k, k \leq N\}$  according to (20), with  $N = 10^7$ ,  $\mu = \mathcal{N}(0, 1)$  and using  $\theta^* = (0.5, 0.7, 0.1)$ . The prior distribution on  $\theta$  is deliberately non-informative and taken as a Gaussian with mean  $(0, 0, 0)$  and a diagonal covariance matrix with a large variance. We restrict the subset parameter to the the set  $\bar{U}_n \subset U_n$  involving  $n$  contiguous observations:

$$\bar{U}_n = \{Y_{0:n-1}, Y_{1:n}, \dots, Y_{N-n+1:N}\} .$$

Using such a sub-window yields a tractable likelihood (21) as otherwise, the function  $h_k$  in (22) is no longer explicit. We compare the efficiency of LWA-MCMC with respect to M-H in function of  $n$ ,  $\epsilon$  and  $S$ . In all this section, the simulations were achieved with a gaussian proposal kernel  $Q$  with variance tuned so that the acceptance rate of the sequence  $\{\theta_k, k \in \mathbb{N}\}$  is between 30% and 40%. For LWA-MCMC, a subset  $U \in \bar{U}_n$  is identified with its starting index and the proposal kernel  $R$  consists in assigning a probability on the discrete alphabet  $\{1, \dots, N_n + 1\}$  standing for all the possible starting time for subsets  $U \in \bar{U}_n$ . More precisely,  $R$  can be written as

$$R(n_0; n'_0) \propto \omega \exp(-\lambda |n_0 - n'_0|) + (1 - \omega) \exp(-\lambda |n' - n'_0|), \\ n' \sim \mathcal{U}(\{1, \dots, N - n + 1\}). \quad (23)$$

The rationale is to propose  $U' \sim R(U, \cdot)$  through a mixture of two distributions: the first gives higher weight to local moves and the latter allows jumps to remote sections of the time series (through the offset  $n'$ ). This mixture allows to browse through  $\bar{U}_n$ , while avoiding to remain trapped in local optimal subsets. In this example, we have used  $\omega = 0.9$  and  $\lambda = 0.1$ .

Figure 4 displays the sample path of three different Markov chains (blue, black and red) simulated during a fixed time budget (1 hour). The blue dashed sample path is the standard M-H targeting the full-posterior  $\pi$  and the other two correspond to the LWA-MCMC transition kernel for  $n = 10,000$  and  $n = 1,000$ , with  $\epsilon = 1$ . In this setup, the summary statistics vector was defined as  $S = S_0$  with

$$\forall U \in \bar{U}_n \quad S_0(Y_U) = (q_{.2}(Y_U), q_{.5}(Y_U), q_{.8}(Y_U), \rho_1(Y_U), \dots, \rho_5(Y_U)) , \quad (24)$$

where for any  $\lambda \in (0, 1)$ ,  $q_\lambda(Y_U)$  is  $Y_U$   $\lambda$ -quantiles and for all  $p \in \mathbb{N}$ ,  $\rho_p(Y_U)$  is the  $p$ -lag sample autocorrelation. As expected, the time to reach stationarity is significantly reduced when sampling through LWA-MCMC. Actually, the M-H sampler does not even reach the high density area of the support during this time frame. On the other hand, when reducing  $n$  the time to reach convergence is even smaller but the stability of the chain may be affected. This is illustrated with Figure 5 which shows the posterior distribution of the three parameters  $\tilde{\pi}_{n,\epsilon}^*$  for  $\epsilon = 1$ ,  $S = S_0$ ,  $n = 10,000$ ,  $n = 1,000$  and  $n = 100$  obtained through a time-normalized LWA-MCMC simulation. For the setup featuring  $n = 100$ , the efficiency of the sampler could be criticized, since a slight bias results. Nonetheless, in an attempt to correct this behaviour, the parameter  $\epsilon$  is decreased, making the choice of subset more critical. Here, Figure 6 shows that it is worthwhile to penalize subsets which are less relevant with respect to the summary statistics  $S_0$ . More precisely, since we are interested in the posterior distribution of the parameter we have reported the mean density observed along the confidence intervals corresponding to the .2-quantiles






Figure 4: ARMA example - M-H (dashed, blue) and LWA-MCMC (dashed, red and black) ( $\epsilon = 1$ ,  $S = S_0$ ) for a fixed computational budget.

and .8-quantiles. Indeed, because each run is a stochastic process (a random collection of subsets is sampled) there is variability within the collection of posterior distributions provided by each run (especially when  $\epsilon \ll 1$ ). These results were obtained through 100 independent runs of LWA-MCMC for  $n = 100$ ,  $S = S_0$  and  $\epsilon \in \{1, 10^{-2}, 10^{-4}\}$ , each run featuring 200,000 iterations among which the 10,000 first were discarded for burn-in. The initial states of all chains were drawn from the prior distribution. Moreover, we have included two extra setups which can be regarded as "extreme" scenarios:

- **free subset** – a standard M-H kernel targeting a sub-posterior which refreshes uniformly the  $n = 100$  data of the subset at each transition,
- **fixed subset** – a standard M-H kernel targeting a sub-posterior with a fixed subset of  $n = 100$  data chosen uniformly before the first iteration.

While the first one can be seen as a LWA-MCMC kernel with  $\epsilon \rightarrow \infty$  since the relevance of the subset does not interfere the parameter sampling, the latter is related with a LWA-MCMC kernel which would be trap in a single subset throughout all the sampling scheme, hence  $\epsilon \rightarrow 0$ . The vertical green and blue lines respectively indicate the true parameter value and the expectation of the mean sub-posterior distribution targeted by the different MCMC. Clearly for  $n = 100$  and  $S = S_0$ , the parameter  $\epsilon = 10^{-2}$  yields an optimal setup since for all the three parameters, the expectation of the posterior distributions meets the true value. Beyond this optimal  $\epsilon$ , the weights  $\{\nu_{n,\epsilon}(U), U \in \bar{U}_n\}$  become too unbalanced. As a result some relevant subsets are simply ignored and the inference is performed on a restricted set of subset. Table 2 shows indeed that for  $\epsilon < 10^{-2}$ , the subset may be refreshed relatively infrequently.

Finally, we consider for the setup  $n = 100$  and  $\epsilon = 10^{-2}$ , two other possible choices

setup / $\epsilon$	free subset	1	$10^{-1}$	$10^{-2}$	$10^{-3}$	$10^{-4}$	fixed subset
$U$ refresh rate	1	.81	.34	.05	.001	$10^{-4}$	0

Table 2: ARMA example - Refresh probability of the subset  $U$  when targeting  $\pi_{n,\epsilon}$  with  $n = 100$  and  $S = S_0$ .




Figure 5: ARMA example - Boxplot of 100 independent LWA-MCMC sample paths for different  $n$  – green lines represent the true parameter  $\theta^*$ .

for  $S$ :

$$S_1(Y_U) = (\rho_1(Y_U), \dots, \rho_{15}(Y_U)) , \quad S_2(Y_U) = (\min(Y_U), \max(Y_U)) .$$

Figure 7 shows that this *naive* choice of  $S_2$  as summary statistics yields a poor sub-posterior target. Indeed, the minimum and the maximum of the time series does not characterize much the model parameters. Note that this setup gives worse results than the free subset setup: while the latter does not advantage any subset  $U \in \mathcal{U}_n$ , the LWA-MCMC setup with  $n = 100$ ,  $\epsilon = 10^{-2}$ ,  $S = S_2$  uses more likely subsets which match the min/max statistics of the full dataset and which, as a result, may force the sampler to infer the parameters through those unrepresentative subsets. Indeed, the min/max stopping times are likely to be far from each other as they represent large deviations to the stationary process. Therefore the subsets whose min/max statistics of the  $n = 100$  contiguous time step are close to that of the full dataset can reasonably be labeled as "anomalies". On the other hand, when the inference is performed using  $S = S_1$ , the collection of subsets is only guided through the correlation between consecutive states. We see that the results are better when these relative statistics are coupled with global statistics such as the mean and the two quantiles  $q_2$  and  $q_8$ , like in  $S = S_0$ .

## 5.2. Binary classification

In this application, we consider the binary classification example used in the MCMC SubLikelihood (MCMCSubLhd) approach (Bardenet et al., 2014). MHSUBLhd is an alternative to M-H, designed to perform Bayesian inference in large datasets contexts; see (2) and the corresponding introduction section for more details. We consider here the example in Section 4.2 of Bardenet et al. (2014) in order to compare the efficiency of LWA-MCMC and MHSUBLhd. A large number ( $N = 10^7$ ) of realizations from a two-dimensional Gaussian mixture distribution with two classes were sampled with the following parameters

$$\mu_1 = [-1, 0], \mu_2 = [1, 0], \Sigma_1 = \Sigma(\sigma_2 = .25), \Sigma_2 = \Sigma(\sigma_2 = .25), \Sigma(\sigma) = \text{diag}(\sigma^2; \sigma^2/2),$$

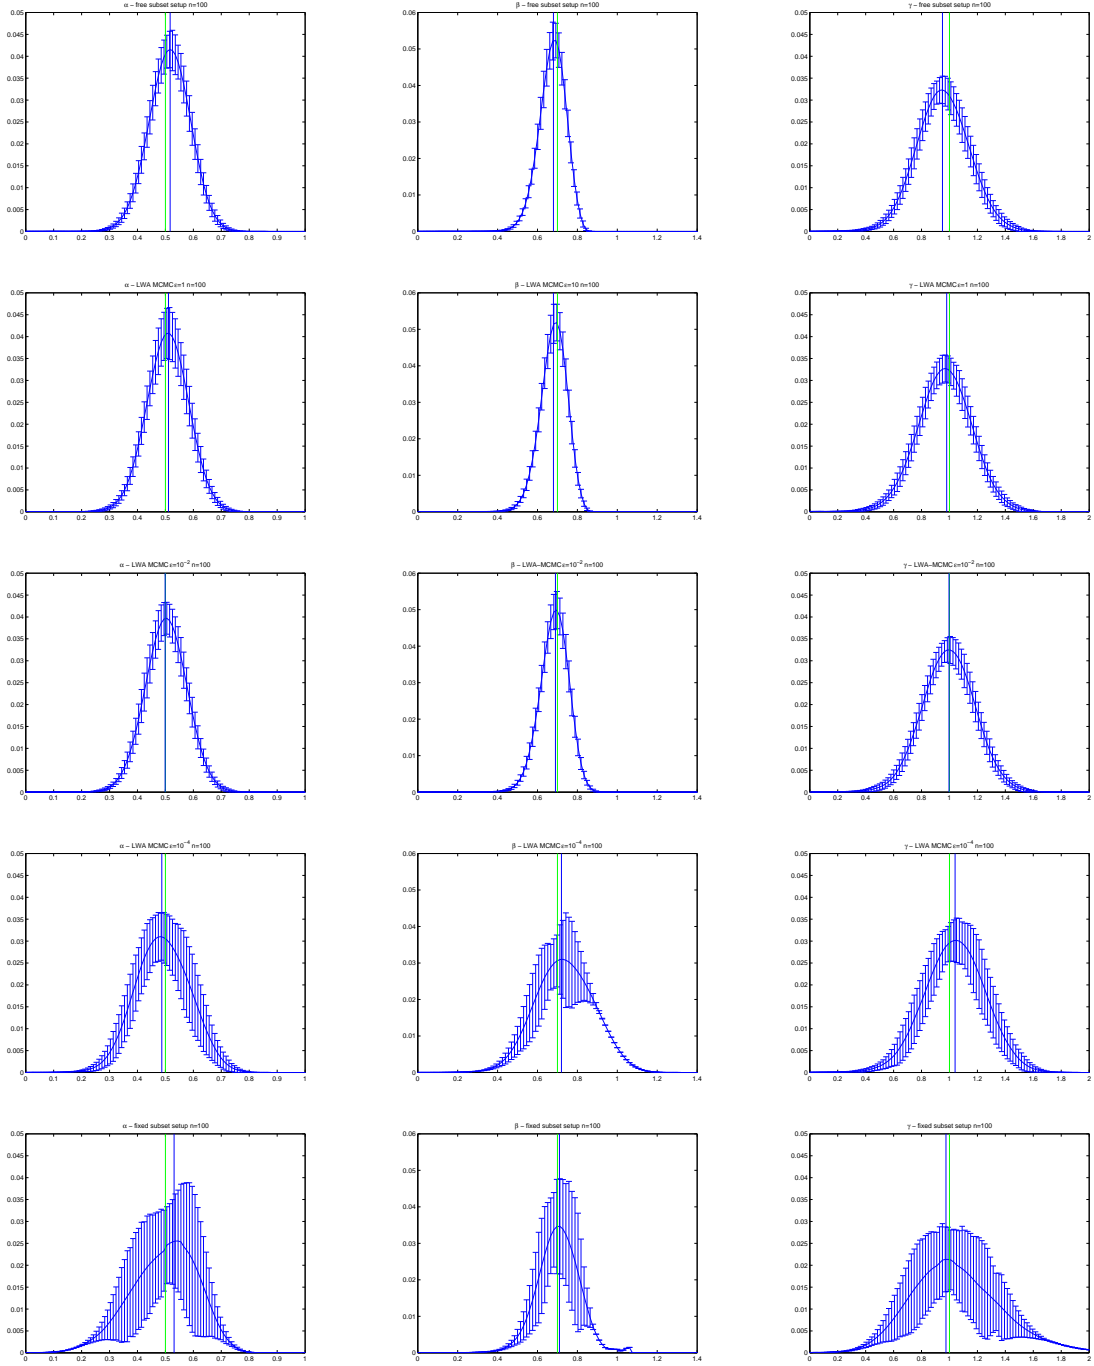


Figure 6: ARMA example - Mean density and confident intervals for each parameter, obtained through 100 independent LWA-MCMC runs and for different  $\epsilon$ . Vertical blue lines give the mean of the mean density and vertical green lines the true parameter  $\theta^*$ .




Figure 7: ARMA example - Mean density and confident intervals for each parameter, obtained through 100 independent LWA-MCMC runs and for different summary statistics  $S$ . Vertical blue lines give the mean of the mean density and vertical green lines the true parameter  $\theta^*$ .




Figure 8: Binary Classification example - Samples from the two-dimensional mixture of two Gaussian distributions.

for the mean and covariance matrix of the two components respectively; see Figure 8. Both LWA-MCMC and MHSubLhd provide a sequence of parameters  $\{\theta_{j,k} = (\mu_{j,k}, \sigma_{j,k}), j \in \{1, 2\}, k \in \mathbb{N}\}$  which is used to classify test-data  $\{\tilde{Y}_m, m \leq 10^7\}$ , sampled from the same model, through the following real-time maximum likelihood classifier defined at time  $t > 0$  by:

$$C_t(\tilde{Y}_m) = \arg \max_{j \in \{1, 2\}} f(Y_u | \theta_{j, \kappa(t)}) \quad \text{where} \quad \begin{cases} \forall t \in \mathbb{R}, \kappa(t) = \max_{k \in \mathbb{N}} \{t \geq \tau_k\}, \\ \tau_k \text{ is the time at the end of the } k\text{-th iteration.} \end{cases}$$

On the basis of the time series simulation example, LWA-MCMC was tuned with  $n = 1,000$ ,  $\epsilon = .01$  and  $nS(U) = [\sum_{k \in U} \mathbb{1}_{\{I_k=0\}}; \sum_{k \in U} \mathbb{1}_{\{I_k=1\}}]$ . The specific MHSubLhd parameters were chosen accordingly:

- $\delta = .1$ , which states that the decision to accept/reject a proposed candidate in MHSubLhd is in accordance with the M-H decision with probability  $1 - \delta = .9$
- every time a decision (accept/reject a candidate) is postponed,  $n_\ell = 1000\ell$  new data are added to the current subset, where  $\ell \in \mathbb{N}^*$  is the number of times that a decision has been postponed in the current MCMC transition (the size of the subset increments, following the guidelines provided in the Section 2.2 of Bardenet et al. (2014)).

scenario	1	2	3	4
LWA-MCMC	1,000	1,000	1,000	1,000
MHSubLhd	$2.4 \cdot 10^6$	$1.9 \cdot 10^6$	$2.35 \cdot 10^6$	$2.8 \cdot 10^6$

Table 3: Binary Classification example - Average number of data used per transition.

Finally, the same proposal kernel were used for both samplers:

$$j \sim \text{Unif}(1, 2), \quad (\epsilon_1, \epsilon_2) \sim \mathcal{N}(0, 1), \quad \begin{cases} \mu'_j = \mu_j + \varsigma_j \epsilon_1 \\ \sigma'_j = \exp(\varrho_j \epsilon_2) \sigma_j \end{cases},$$

where the parameter  $\{\varsigma_j, \varrho_j\}_{j \in \{1, 2\}}$  were updated to maintain an acceptance rate between .25 and .35 through an adaptive Metropolis procedure (Haario et al., 2001).

Figure 9 shows four independent comparisons between LWA-MCMC and MHSubLhd. Both samplers starts with the same initial state (drawn from the prior) in all scenarios. The first column shows the sample path of the two Markov chains against the time (in second): plain lines are LWA-MCMC sample paths and dashed lines are MHSubLhd sample paths. We stress that the step shape of the MHSubLhd sample paths does not highlights a poor mixing chain but illustrates the fact that a single MHSubLhd transition can take a similar amount of time as a standard M-H transition. As a consequence, the chain remains at the same state for a large amount of time. Table 3 shows indeed that, on average, the MHSubLhd sampler ends up using almost 25% of the full dataset at each transition. The results clearly show that, in this application, using LWA-MCMC instead of MHSubLhd results in a practically useful approach, with some spectacular convergence acceleration: in scenario 4, LWA-MCMC is 22 times faster than MHSubLhd.

### 5.3. Additional details for the handwritten digit example of Section 2

In the handwritten digit example (see Section 2), we have used batches of  $n = 100$  data. The summary statistics were simply defined so as to promote subsets which have 20 observations from each class. A very low bandwidth  $\epsilon = 10^{-5}$  was used in order to enforce this characteristic. Similarly to the binary classification example, the proposal kernels of LWA-MCMC and M-H were defined in the same style, namely a Random Walk kernel where at each iteration only a bloc of the template parameter of one of the 5 classes is updated ( $\theta \in \mathbb{R}^{256}$  in this example). The variance parameter of the Random Walk is here again adapted according to the past trajectory of the chain, so as to maintain an acceptance rate of .25.

In contrast to the previous examples, the computational difference is not explained by the fact that the M-H acceptance probability is more expensive to compute. Indeed, both need to evaluate the function  $\phi(\theta'_j) : \mathbb{R}^{256} \rightarrow \mathbb{R}^{225}$  for the updated class  $j$ , which represents the only heavy routine calculation. In fact, the adjusted variance of the M-H proposal turns out to be 10 times lower than that for LWA-MCMC. Loosely speaking, on the one hand, the M-H proposal needs to provide a parameter  $\theta'$  which fits in about 2,000 observations: the log of the acceptance ratio depends of  $\sum_{k=1}^N \mathbb{1}_{I_k=j} \{\|Y_k - \phi(\theta'_j)\|^2 - \|Y_k - \phi(\theta_j)\|^2\}$  for the updated class  $j$ . On the other hand, the LWA-MCMC proposal should match only about 20 images through the quadratic term  $\sum_{k \in U} \mathbb{1}_{I_k=j} \{\|Y_k - \phi(\theta'_j)\|^2 - \|Y_k - \phi(\theta_j)\|^2\}$ . As a consequence, the M-H adapted variance makes the Random Walk

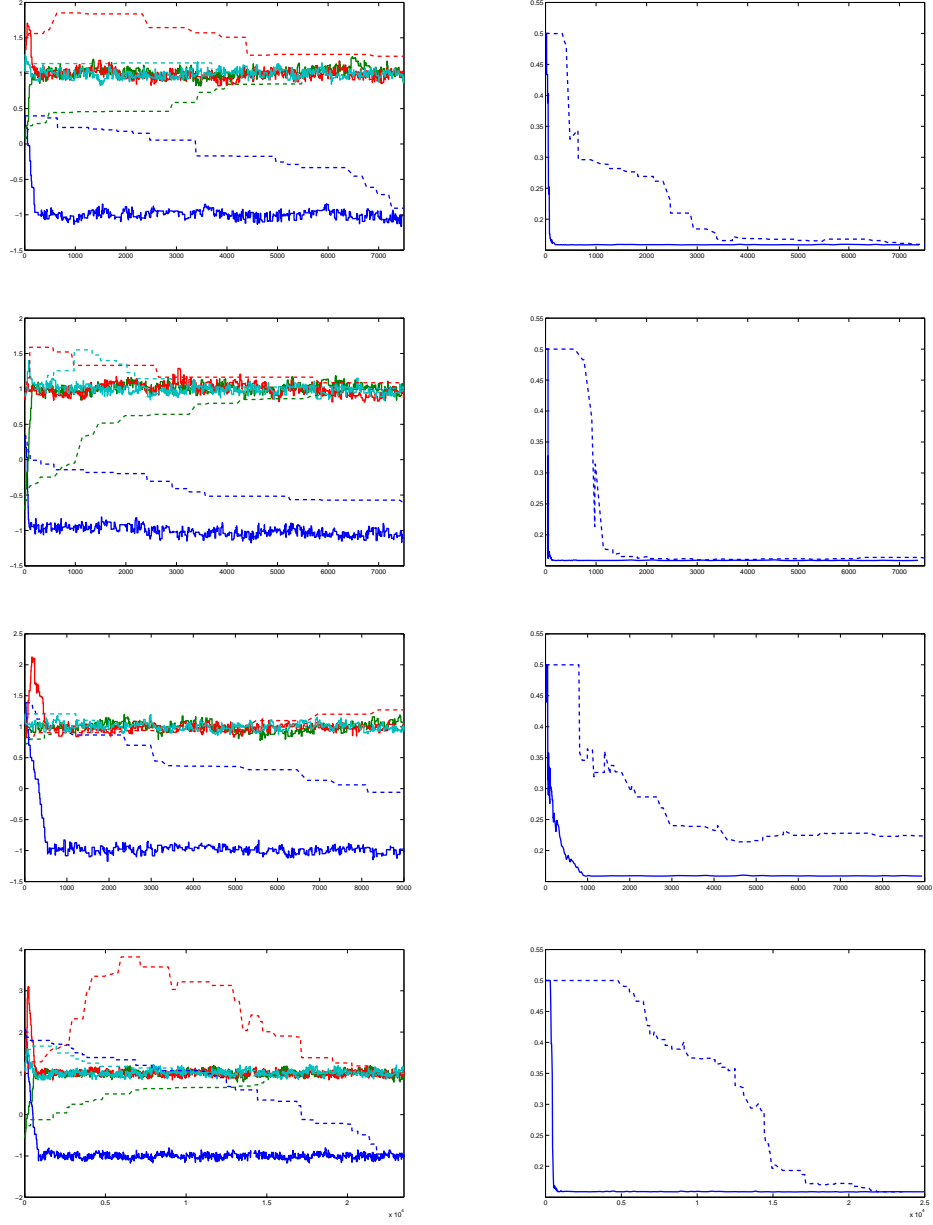


Figure 9: Binary Classification example - Four independent scenarios – left column: sample paths of the two samplers (LWA-MCMC in plain and MHSubLhd dashed) ; right column: classification error rate on the training dataset (LWA-MCMC in plain and MHSubLhd dashed).

less efficient, which results in the Markov chain  $\{\theta_k, k \in \mathbb{N}\}$  exploring the state space  $\Theta$  slower.

## 6. Conclusion

The Light and Widely Applicable MCMC methodology introduced and discussed in this paper is attractive as it overcomes the critical issues encountered by the popular Metropolis–Hastings sampler in the modern development of big data inference problems. Several recent noisy M–H methods have been proposed to address these issues. However, (i) they are only valid for *i.i.d.* realizations and (ii) they may use a significant portion of the dataset negating any potential computational saving. LWA–MCMC pushes the approximation one step forward to preserve the celebrated M–H simplicity. The efficiency of the sampler is illustrated in several typical Bayesian problems including parameter estimation and classification.

Our experiments have illustrated the usefulness of LWA–MCMC. Future work should extend the theory beyond the scope of exponential models.

## Acknowledgements

Florian Maire thanks The Insight Centre for Data Analytics for funding the Post-Doctoral Fellowship. The Insight Centre for Data Analytics is supported by Science Foundation Ireland under Grant Number SFI/12/RC/2289. Nial Friel’s research was also supported by an Science Foundation Ireland grant: 12/IP/1424.

## Appendix A. Proof of Proposition 1

*Proof.* Without loss of generality, we take  $g$  as the identity on  $\Theta$ . By straightforward algebra,

$$\begin{aligned} \text{KL}(\pi \| \tilde{\pi}_U) &= \mathbb{E}_\pi \left\{ \log \frac{\pi(\theta)}{\tilde{\pi}_U(\theta)} \right\} = \left\langle \mathbb{E}_\pi(\theta), \sum_{k=1}^N S(Y_k) - \sum_{k \in U} S(Y_k) \right\rangle + (n - N) \mathbb{E}_\pi \{ \ell(\theta) \} , \\ &\quad + \log \frac{Z(Y_U)}{Z(Y_{1:N})} , \quad (\text{A.1}) \end{aligned}$$

where  $\ell(\theta) = \log L(\theta)$ . Now, let

$$\xi_U = \frac{1}{n} \sum_{k \in U} S(Y_k) - \frac{1}{N} \sum_{k=1}^N S(Y_k) \quad (\text{A.2})$$

and express (A.1) as:

$$\begin{aligned} \text{KL}(\pi \| \tilde{\pi}_U) &= \left(1 - \frac{n}{N}\right) \left\langle \mathbb{E}_\pi(\theta), \sum_{k=1}^N S(Y_k) \right\rangle + (n - N) \mathbb{E}_\pi \{ \ell(\theta) \} - n \langle \xi_U, \mathbb{E}_\pi(\theta) \rangle , \\ &\quad + \log \frac{Z(Y_U)}{Z(Y_{1:N})} . \quad (\text{A.3}) \end{aligned}$$

We now want to find an upper bound of  $\log Z(Y_U)/Z(Y_{1:N})$ . First, note that

$$\begin{aligned} Z(Y_U) &= \int p(d\theta) \frac{\exp(n/N) \left\langle \theta, \sum_{k=1}^N S(Y_k) \right\rangle}{L(\theta)^n} \exp(n \langle \theta, \xi_U \rangle) \\ &= \int p(d\theta) \left\{ \prod_{k=1}^N f(Y_k | \theta) \right\}^{\frac{n}{N}} \exp(n \langle \theta, \xi_U \rangle) \\ &= \int p(d\theta) \prod_{k=1}^N f(Y_k | \theta) \exp(n \langle \theta, \xi_U \rangle) \left\{ \prod_{k=1}^N f(Y_k | \theta) \right\}^{\frac{n}{N}-1} \quad (\text{A.4}) \end{aligned}$$

Because  $n/N \in (0, 1)$ , we can write

$$\left\{ \prod_{k=1}^N f(Y_k | \theta) \right\}^{\frac{n}{N}-1} \leq \alpha(Y_{1:n})^{\frac{n}{N}-1}$$

where  $\alpha(Y_{1:n}) = \inf_{\theta \in \Theta} \prod_{k=1}^N f(Y_k | \theta)$ . As a consequence, we have:

$$\begin{aligned} \frac{Z(Y_U)}{Z(Y_{1:N})} &\leq \alpha(Y_{1:n})^{\frac{n}{N}-1} \int \pi(d\theta | Y_{1:N}) \exp(n \langle \theta, \xi_U \rangle) \\ &= \alpha(Y_{1:n})^{\frac{n}{N}-1} \mathbb{E}_\pi \{ \exp(n \langle \theta, \xi_U \rangle) \} \quad (\text{A.5}) \end{aligned}$$



Plugging (A.5) into (A.3) yields

$$\text{KL}(\pi \| \tilde{\pi}_U) \leq \underbrace{\left(1 - \frac{n}{N}\right) \left\{ \left\langle \mathbb{E}_\pi(\theta), \sum_{k=1}^N S(Y_k) \right\rangle - \log \alpha(Y_{1:N}) \right\}}_{\Psi(n, N, Y_{1:N})} + (n - N) \mathbb{E}_\pi \{ \ell(\theta) \} \\ + \underbrace{\log \mathbb{E}_\pi \{ \exp(n \langle \theta, \xi_U \rangle) \} - n \langle \xi_U, \mathbb{E}_\pi(\theta) \rangle}_{\Omega(U, n, Y_{1:N})}. \quad (\text{A.6})$$

Now, using the Cauchy-Schwartz inequality we have:

$$\Omega(U, n, Y_{1:N}) = \log \mathbb{E}_\pi [\exp \{n \langle \theta, \xi_U \rangle - n \langle \xi_U, \mathbb{E}_\pi(\theta) \rangle\}] \\ = \log \mathbb{E}_\pi [\exp \{n \langle \theta - \mathbb{E}_\pi(\theta), \xi_U \rangle\}] \leq \underbrace{\log \mathbb{E}_\pi \{ \exp(n \|\theta - \mathbb{E}_\pi(\theta)\| \cdot \|\xi_U\|) \}}_{B(U, n, Y_{1:N})}.$$

Finally we have

$$\text{KL}(\pi \| \tilde{\pi}_U) \leq \Psi(n, N, Y_{1:N}) + B(U, n, Y_{1:N}), \quad \text{where } B(U, n, Y_{1:N}) \geq 0$$

such that the subset  $U \in \mathcal{U}_n$  minimizing  $\|\xi_U\|$  also minimize  $B(U, n, Y_{1:N})$  with the special case

$$U \in \mathcal{U}_n^* \Rightarrow B(U, n, Y_{1:N}) = 0,$$

hence completing the proof.  $\square$

## References

- Allasonnière, S., Amit, Y., Trouvé, A., 2007. Towards a coherent statistical framework for dense deformable template estimation. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 69 (1), 3–29.
- Alquier, P., Friel, N., Everitt, R., Boland, A., 2014. Noisy Monte Carlo: Convergence of Markov chains with approximate transition kernels. *Statistics and Computing*, to appear.
- Andrieu, C., Roberts, G. O., 2009. The pseudo-marginal approach for efficient Monte Carlo computations. *The Annals of Statistics*, 697–725.
- Andrieu, C., Vihola, M., 2015. Convergence properties of pseudo-marginal markov chain monte carlo algorithms. *The Annals of Applied Probability* 25 (2), 1030–1077.
- Banterle, M., Grazian, C., Robert, C. P., 2014. Accelerating Metropolis-Hastings algorithms: Delayed acceptance with prefetching. *arXiv preprint arXiv:1406.2660*.
- Bardenet, R., Doucet, A., Holmes, C., 2014. Towards scaling up Markov chain Monte Carlo: an adaptive subsampling approach. In: *Proceedings of the 31st International Conference on Machine Learning*. pp. 405–413.
- Csilléry, K., Blum, M. G., Gaggiotti, O. E., François, O., 2010. Approximate Bayesian computation (ABC) in practice. *Trends in ecology & evolution* 25 (7), 410–418.
- Fearnhead, P., Prangle, D., 2012. Constructing summary statistics for approximate Bayesian computation: semi-automatic approximate Bayesian computation. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 74 (3), 419–474.
- Haario, H., Saksman, E., Tamminen, J., 2001. An adaptive Metropolis algorithm. *Bernoulli*, 223–242.
- Korattikara, A., Chen, Y., Welling, M., 2014. Austerity in MCMC land: Cutting the Metropolis-Hastings budget. In: *Proceedings of the 31st International Conference on Machine Learning*.
- Kuhn, E., Lavielle, M., 2004. Coupling a stochastic approximation version of EM with an MCMC procedure. *ESAIM: Probability and Statistics* 8, 115–131.
- Maclaurin, D., Adams, R. P., 2014. Firefly monte carlo: Exact MCMC with subsets of data. *arXiv preprint arXiv:1403.5693*.

- Marin, J.-M., Pudlo, P., Robert, C. P., Ryder, R. J., 2012. Approximate Bayesian computational methods. *Statistics and Computing* 22 (6), 1167–1180.
- Nunes, M. A., Balding, D. J., 2010. On optimal selection of summary statistics for approximate Bayesian computation. *Statistical applications in genetics and molecular biology* 9 (1).
- Pillai, N. S., Smith, A., 2014. Ergodicity of approximate mcmc chains with applications to large data sets. *arXiv preprint arXiv:1405.0182*.
- Robbins, H., Monro, S., 1951. A stochastic approximation method. *The annals of mathematical statistics*, 400–407.
- Tierney, L., 1998. A note on Metropolis-Hastings kernels for general state spaces. *Annals of Applied Probability*, 1–9.
- Van der Vaart, A. W., 2000. *Asymptotic statistics*. Vol. 3. Cambridge university press.